arXiv:1511.05987v1 [cs.DS] 18 Nov 2015

Algorithms for Communication Problems for Mobile Agents Exchanging Energy

Jerzy Czyzowicz[2⋆], Krzysztof Diks[1⋆⋆], Jean Moussi[2], and Wojciech Rytter[1⋆⋆⋆]

[1] Faculty of Mathematics, Informatics and Mechanics,
University of Warsaw, Warsaw, Poland [diks,rytter]@mimuw.edu.pl
[jurek,Jean.Moussi]@uqo.ca
[2] Département d'informatique, Université du Québec
en Outaouais, Gatineau, Québec, Canada

Abstract. We consider communication problems in the setting of mobile agents deployed in an edge-weighted network. The assumption of the paper is that each agent has some energy that it can transfer to any other agent when they meet (together with the information it holds). The paper deals with three communication problems: data delivery, convergecast and broadcast. These problems are posed for a centralized scheduler which has full knowledge of the instance. It is already known that, without energy exchange, all three problems are NP-complete even if the network is a line. Surprisingly, if we allow the agents to exchange energy, we show that all three problems are polynomially solvable on trees and have linear time algorithms on the line. On the other hand for general undirected and directed graphs we show that these problems, even if energy exchange is allowed, are still NP-complete.

1 Introduction

A set of n agents is placed at nodes of an edge-weighted graph G. An edge weight represents its length, i.e., the distance between its endpoints along the edge. Each agent has an amount of energy (possibly distinct for different agents). Agents walk in a continuous way along the network edges using amount of energy proportional to the distance travelled.

An agent may stop at any point of a network edge (i.e. at any distance from the edge endpoints, up to the edge weight). Each agent has memory in which it can store information.

When two agents meet, one of them can transfer a portion of currently possessed energy to another one. Moreover, two meeting agents exchange their currently possessed information, so that after the meeting each of them keeps in its memory the union of pieces of information previously hold by each of them.

We assume that each agent has sufficient memory to store the entire information initially belonging to all agents.
Our algorithms work as centralized schedulers having full knowledge of the instance.

⋆ Supported by NSERC grant of Canada
⋆⋆ Supported by the grant NCN2014/13/B/ST6/00770 of the Polish Science Center.
⋆⋆⋆ Supported by the grant NCN2014/13/B/ST6/00770 of the Polish Science Center.

We consider three problems:

1. ***Data delivery* problem:** Given two nodes s, t of G, is it possible to transfer the initial packet of information placed at node s to node t?
2. ***Convergecast* problem:** Is it possible to transfer the initial information possessed by each agent to a fixed agent ? (See Fig.1)
3. ***Broadcast* problem:** Is it possible to transfer the initial information of some agent to all other agents ? (See Fig.1)

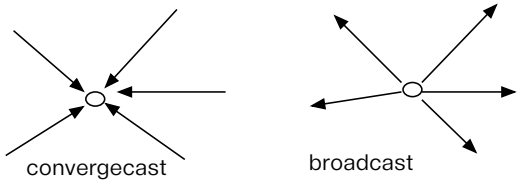

Fig. 1. Schematic view of convergecast and broadcast.

We will look for schedules of agent movements which will not only result in completing the desired task, but also attempt to maximize the unused energy. We call such schedules *optimal*. We conservatively suppose that, whenever two agents meet, they automatically exchange the entire information they held. This information exchange procedure is never explicitly mentioned in our algorithms, supposing, by default, that it always takes place when a meeting occurs.

1.1 Our results

We show that all three communication problems are polynomially solvable on trees and have linear time algorithms on the line. On the other hand for general undirected and directed graphs we show that these problems, even if energy exchange is allowed, are still NP-complete.

1.2 Related work

Recent development in the network industry triggered the research interest in mobile agents computing. Several applications involve physical mobile devices like robots, motor vehicles or various wireless gadgets. Mobile agents are sometimes interpreted as software agents, i.e., programs migrating from host to host in a network, performing some specific tasks. Examples of agents also include living beings: humans (e.g. soldiers or disaster relief personnel) or animals. Most studied problems for mobile agents involve some sort of environment search or exploration (cf. [3,9,12,13]). In the case of a team of collaborating mobile agents, the challenge is to balance the workload among the agents in order to minimize the exploration time. However this task is often hard (cf. [14]), even in the case of two agents in a tree, [7].

The task of convergecast is important when agents possess partial information about the network (e.g. when individual agents hold measurements performed by sensors located at their positions) and the aggregate information is needed to make some global decision based on all measurements. The convergecast problem is often considered as a version of the data aggregation question (e.g. [17,18]) and it has been investigated mainly in the context of wireless and sensor networks, where the energy consumption is an important issue (cf. [5,16]).

The task of broadcast is useful, e.g., when a designated leader needs to share its information with collaborating agents in order to perform together some future tasks.

The broadcast problem for stationary processors has been extensively studied both in the case of the message passing model, (e.g. [4]), and for the wireless model, (see [10]).

The power awareness question has been studied in different contexts. Energy management of (not necessarily mobile) computational devices has been studied in [2]. To reduce energy consumption of computer systems the methods proposed include power-down strategies (see [2,6,15]) or speed scaling (cf. [19]). Most of research on energy efficiency considers optimization of overall power used. When the power assignments are made by the individual system components, similar to our setting, the optimization problem involved has a flavor of load balancing (cf. [8]).

The broadcast problem for stationary processors has been extensively studied both in the case of the message passing model, (e.g. [4]), and for the wireless model, (see [10]).

The problem of communication by energy-constrained mobile agents has been investigated in [1]. The agents of [1] all have the same initial energy and they perform efficient convergecast and broadcast in line networks. However the same problem for tree networks is proven to be strongly NP-complete in [1].

The closely related problem of data delivery, when the information has to be transmitted between two given network nodes by a set of energy constrained agents has been studied in [11]. This problem is proven to be NP-complete in [11] already for line networks, if the initial energy values may be distinct for different agents. However, in the setting studied in [1,11], the agents do not exchange energy. In the present paper we show that the situation is quite different if the agents are allowed to transfer energy between one another.

2 The line environment

In this section we suppose that we are given a collection of agents $\{0, 1, 2, \ldots, n-1\}$ on the line. Each agent i is initially placed at position a_i on the line and has initial energy $e(i)$. We investgate delivery, convergecast and broadcast problems separately. The are solved using auxiliary table.s

2.1 Data delivery on the line

We start with the delivery problem from point s to t. Assume w.l.o.g. that $a_i < a_j$ for $i < j$ and $s < t$.

The problem can be immediately reduced to the situation $s = a_1,\ t = a_n$. Otherwise the first agent is going to s from left to right, swallawing energy of encountered agents, symetrically the righmost agent is going right to left until reaching t. Then we can reduce n and renumber agents setting $a_1 = s,\ a_n = t$.

Our first algorithm is only a decision version. Its main purpose is to show how certain useful table can be computed, all other algorithms are based on computing similar type of tables.

Consider the partial delivery problem $\mathcal{D}_i$ which is the original problem with agents larger than i removed, together with their energy, and the goal is to deliver the packet from the 1-st agent to the i-th agent.

We say that the problem $\mathcal{D}_i$ is *solvable* iff such a delivery is possible.

We define the following table $\overrightarrow{\Delta}$:

- If $\mathcal{D}_i$ is not solvable then $\overrightarrow{\Delta}(i) = -\delta$, where δ is the *minimal* energy which needs to be added to e_i (to energy of i-th agent) to make $\mathcal{D}_i$ solvable.
- If $\mathcal{D}_i$ is solvable then $\overrightarrow{\Delta}(i)$ is the *maximal* energy which can remain in point a_i after delivering the packet from a_1 to a_i. Possibly $\overrightarrow{\Delta}(i) > e(i)$ since during delivery the partial energy of some other agents can be moved to point a_i.

Assume in the algorithm that points s and t are the starting points $s = a_1$ and $t = a_n$. In our algorithm the statements of the form $x+ = y$ are equiavalnet to $x := x + y$.

ALGORITHM DELIVERY-TEST-ON-THE-LINE ;

{ Decision version and computation of table $\overrightarrow{\Delta}$ }

1. **for each** $i \in [0..n-1]$ **do** $\overrightarrow{\Delta}(i) := e(i)$;
2. **for** $i = 1$ to n **do**
3. $\quad d := a_i - a_{i-1}$;
4. $\quad$ **if** $\overrightarrow{\Delta}(i-1) \geq d$ **then** $\overrightarrow{\Delta}(i) \mathrel{+}= \overrightarrow{\Delta}(i-1) - d$
5. $\quad$ **else if** $\overrightarrow{\Delta}(i-1) \geq 0$ **then** $\overrightarrow{\Delta}(i) \mathrel{+}= 2(\overrightarrow{\Delta}(i-1) - d)$
6. $\quad$ **else** $\overrightarrow{\Delta}(i) \mathrel{+}= \overrightarrow{\Delta}(i-1) - 2d$;
8. Delivery from a_1 to a_n is possible iff $\overrightarrow{\Delta}(n) \geq 0$

Example 1. Assume $[a_1, a_2, \ldots a_5] = [0, 10, 20, 30, 40, 50]$, and $[e(1), e(2), \ldots e(5)] = [0, 24, 10, 40, 0]$. Then $\overrightarrow{\Delta} = [0, 4, -2, 18, 8]$.

Remark. The values of $\overrightarrow{\Delta}(i)$ are not needed to solve the decision-only version. However they will be useful in creating the delivery schedule and also in the *convergecast* problem.

Lemma 2. *The algorithm* DELIVERY-TEST-ON-THE-LINE *correctly computes the table* $\overrightarrow{\Delta}$ *(thus it solves the decision version of the delivery problem) in linear time.*

Proof. Asume the algorithm computed correctly $\overrightarrow{\Delta}(i-1]$. There are three cases:

Case 1 (statement 4).

The instance $\mathcal{D}_{i-1}$ is solvable and after moving the packet from a_1 to a_{i-1} the maximal remaining energy is $\overrightarrow{\Delta}(i-1]$. In this case this energy is sufficient to move the packet from a_{i-1} to a_i, together with remaining energy. We spent d energy to travers the distance d in one direction. Then we get in total $\mathcal{D}_{i-1} + \mathcal{D}_i - d$ energy in a_i.

Case 2 (statement 5).

The instance $\mathcal{D}_{i-1}$ is still solvable but after moving the packet from a_1 to a_{i-1} the remaining energy is not sufficient to reach a_i without *help* from agents to the right of a_{i-1}. Then the $(i-1)$-st agent moves only one-way by distance $\overrightarrow{\Delta}(i-1]$. The remaining distance (to a_i) should be covered both-ways from a_i. Hence a_i looses $2(d - \overrightarrow{\Delta}(i-1])$ energy, which is expressed by statement 5.

Case 3 (statement 6).

The instance $\mathcal{D}_{i-1}$ is not solvable, now the node a_{i-1} needs $|\overrightarrow{\Delta}(i-1]|$ additional energy which has to be delivered from a_i. Additionally the i-th agent needs $2d$ energy to cover the distance d from a_i to a_{i-1} in both directions. Consequently its energy is reduced by $2d + |\overrightarrow{\Delta}(i-1]|$. This is reflected in the statement 6.

The cases correspond to the statements in the algorithm, and show their correctness. This completes the proof.

Once the values of $\overrightarrow{\Delta}(i)$ are computed, the schedule describing the behavior of each agent is implicitly obvious, but we give it above for reference. Note that the action of each agent a is started once the process involving lower-numbered agents has been completed. We are not interested in this paper in finding the time to complete the schedule allowing agents to work in parallel.

Observation. In the scheduling algorithm if an agent has the packet (in particular if it is the first agent) then it is not going left, since it holds the packet already.

ALGORITHM DELIVERY-SCHEDULE-ON-THE-LINE ;

{ Delivering packet from a_1 to a_n }

$pos := a_1$;

for $i = 1$ to n **do**

if $\overrightarrow{\Delta}(a) \geq 0$ and $pos \leq a_i$ **then**

1. The i-th agent walks left swallowing energy of encountered agents until arriving at the packet position. It picks up the packet.

2. The i-th agent walks right swallowing energy of encountered agents until exhausting his energy or reaching a_n.

3. Leave the packet at th actual position pos of the i-th agent.

Delivery is successful iff $pos = a_n$;

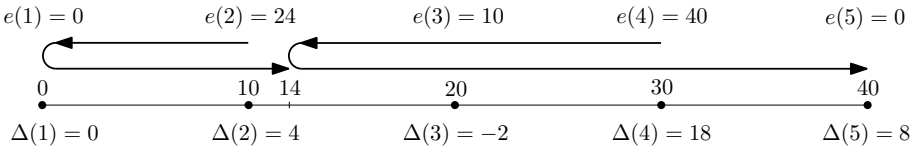

Fig. 2. Schedule of agent movements for a_i's and energies given in Example 1.

We conclude with the following theorem.

Theorem 3. *In linear time we can decide if the information of any agent can be delivered to any other agent and if it is possible find the centrelized scheduling algorithm which perorms such a delivery.*

In the next section we show that the solution may be extended to the convergecast problem.

2.2 The convergecast on the line

The convergecast consists in communication in which the union of initial information of all agents is in the hands of the same agent. Sometimes the problem consists in verifying whether a particular agent is a convergecast agent and other times it has to be determined whether any such agent exists. For energy exchanging agents, if convergecast is possible any agent may be its target, as agents may swap freely when meeting.

We say that a schedule results in convergecast iff there exists an agent c, for $1 \leq c \leq n$ (called a *convergecast agent*), and a schedule permitting to collect by agent c the initial information of any other agent.

ALGORITHM CONVERGECAST-ON-THE-LINE;

1. Compute the values of $\overrightarrow{\Delta}(i)$ and $\overleftarrow{\Delta}(i)$ representing the energy potentials at a_i for deliveries from a_1 to a_n and a_n to a_1, respectively
2. **for** $i = 1$ to n **do**
3. $\quad d := a_{i+1} - a_i$;
4. $\quad$ **if** $(\overrightarrow{\Delta}(i) \geq 0) \wedge (\overrightarrow{\mathcal{B}}(i+1) \geq 0)$ **then** $E := \overrightarrow{\Delta}(i) + \overleftarrow{\Delta}(i) - d$;
5. $\quad\quad (\overrightarrow{\Delta}(i) \geq 0) \wedge (\overleftarrow{\Delta}(i) < 0)$ **then** $E := \overrightarrow{\Delta}(i) + \overleftarrow{\Delta}(i)/2 - d$;
6. $\quad\quad (\overrightarrow{\Delta}(i) < 0) \wedge (\overrightarrow{\mathcal{B}}(i+1) \geq 0)$ **then** $E := \overrightarrow{\Delta}(i)/2 + \overleftarrow{\Delta}(i) - d$;
7. $\quad$ **if** $E \geq 0$ **then return** Convergcast possible;
8. **return** Convergcast not possible;

We have the following theorem.

Theorem 4. *Algorithm* CONVERGECAST-ON-THE-LINE *in* $O(n)$ *time solves the convergecast problem.*

2.3 The broadcast on the line

In the broadcast communication one agent has to transfer its original information to all other agents of the collection. In this section we present an algorithm determining which agents are able to broadcast. Contrary to the convergecast problem, for the energy transferring agents, only a selected subset of them may be able to perform broadcast. However the approach used for convergecast may be transformed the way it is also useful for broadcast communication. We design the schedule, which not only performs broadcast, whenever possible, but also tries to use as little energy of the broadcasting agent as possible. We call such schedule *optimal*.

Our algorithm will compute for all agents the values of $\overrightarrow{\mathcal{B}}(i)$ and $\overleftarrow{\mathcal{B}}(i)$. The value of $\overrightarrow{\mathcal{B}}(i)$ equals the potential of energy at point a_i for the delivery of initial information of agent i to agent 1 using only agents $1, 2, \ldots, i$. More exactly, if $\overrightarrow{\mathcal{B}}(i) < 0$, then $-\Delta(i)$ equals the minimal amount energy which must be added to $e(i)$ so that the delivery from a_i to a_1 is possible. may be suppressed from $e(i)$ so that the delivery from a_i to a_1 be possible.

The values of $\overleftarrow{\mathcal{B}}(i)$ are defined symmetrically.

The following algorithm may be viewed as a different version solving the delivery problem. However the main reason of its presentation is to compute $\overrightarrow{\mathcal{B}}(i)$ values, which are used for broadcast schedule.

ALGORITHM DELIVERY2-ON-THE-LINE;

1. $e := 0; s := a_1;$
2. **for** $i = 1$ to $n + 1$ **do**
3. **if** $a_i > e(i) + e + s$ **then** $e := e + e(i)$; $\overrightarrow{\mathcal{B}}(i) := s + e - a_i$;
4. **else**
   $s += (e + e(i) + \max(a_i - s, 0))/2$;
   $e := 0$; $\overrightarrow{\mathcal{B}}(i) := 2(s - a_i)$;
5. **return** *true* iff $(s \geq a_n)$;

Lemma 5.
(a) *If* $\overrightarrow{\mathcal{B}}(i) \geq 0$ *in algorithm* DELIVERY2-ON-THE-LINE *then agents* $1, 2, \ldots, i$ *can deliver the packet from agent* i *at initial position* a_i *to all agents* $1, 2, \ldots, i-1$. *Moreover,* $\overrightarrow{\mathcal{B}}(i)$ *equals the maximal amount of energy which may be suppressed from* $e(i)$ *(and* $e(i-1), e(i-2), \ldots, e(1)$*, if possible, in this order) so that this delivery is still possible.*

(b) *The tables* $\overrightarrow{\mathcal{B}}(i)$, $\overleftarrow{\mathcal{B}}(i)$ *can be computed in linear time.*

Proof. We prove by induction that the following assertion is true at the completion of i-th iteration of the "for" loop of the algorithm: *Let* k *be the largest integer* $k \leq i$*, such that agents* $1, 2, 3 \ldots, k$ *can deliver the packet from point* s*, such that* $s \geq a_k$*, to point* a_1 *and* s *is the largest possible point having this property. Then*

1. *either* $k = i$, $s \geq a_i$, $e = 0$ *and* $\overrightarrow{\mathcal{B}}(i) = 2(s - a_i) \geq 0$,
2. *or* $k < i$, $s < a_i$ *and* $\overrightarrow{\mathcal{B}}(i) = s + e - a_i < 0$, *where* $e = \sum_{i=k+1}^{i-1} e(i)$

The assertion is clearly true for $i = 1$ as $a_1 = s, e = 0$ and $\overrightarrow{\mathcal{B}}(1) = 0$. Suppose the assertion was true after iteration $i - 1$ and all previous iterations.

Consider first the case 1, i.e. that $k = i - 1$ agents could deliver to a_1 the packet from point $s \geq a_{i-1}$ and at the completion of the $(i-1)$-th iteration $e = 0$ and $DC[i - 1]s - a_{i-1})$. Suppose first that at the i-th iteration the condition at line 3 is true, i.e. $a_i > e(i) + e + s$. That means that agent 1 cannot deliver the packet from point a_i to s and $\overrightarrow{\mathcal{B}}(i) = s + e - a_i < 0$, so the case 2 of the assertion is true. Suppose now that $a_i \leq e(i) + e + s$ at line 3 of the i-th iteration. The agents $1, 2, \ldots, k$ could deliver the packet from $s \geq k$ to a_1. Observe that e contains the accumulated energy of agents $k + 1, k + 2, \ldots, i - 1$, i.e.

$$e = \sum_{i=k+1}^{i-1} e(i)$$

Therefore the condition $a_i \leq e(i) + e + s$ implies that agents $k + 1, \ldots, i$ can deliver the packet from point a_i to s. Hence agents $1, 2, \ldots, i$ can deliver from a_i to a_1. Since $e = 0$ and $\overrightarrow{\mathcal{B}}(i) = 2(s - a_i) \geq 0$, the case 1 of the assertion becomes

true. Observe that the new point s, computed at line 4 is at distance equal twice the excess of energy obtained at a_i so it is computed correctly.

Take now the case that $i-1$ agents could not deliver to a_1 the packet from point a_{i-1} and at the completion of the $(i-1)$-th iteration we have $s < a_i$ and $\overrightarrow{\mathcal{B}}(i) = s + e - a_i < 0$. Suppose first that at the i-th iteration the condition at line 3 is true, i.e. $a_i > e(i) + e + s$. Then, it means that the packet cannot be delivered from a_i to s using agents $k+1, k+2, \ldots, i-1$ and in line 3 we have $e = \sum_{i=k+1}^{i} e(i)$, so the case 2 of the assertion is true. Suppose now that $a_i \leq e(i) + e + s$ at line 3 of the i-th iteration. Similarly to the case above, agents $1, 2, \ldots, k$ could deliver the packet from $s \geq a_k$ to a_1 and agents $k+1, \ldots, i$ can deliver the packet from point a_i to s and the case 1 of the assertion becomes true. This completes the proof.

ALGORITHM Broadcast-on-the-Line;

1. $BR := \emptyset$;
2. Compute the values of $\overrightarrow{\mathcal{B}}(i)$ and $\overleftarrow{\mathcal{B}}(i)$ representing the energy potential for deliveries from a_1 to a_n and a_n to a_1, respectively
3. **for** $i = 1$ to n **do**
4. **if** $\overrightarrow{\mathcal{B}}(i) + \overleftarrow{\mathcal{B}}(i+1) - 2(a_{i+1} - a_i) \geq 0$ **then**
5. $BR := BR \cup \{i, i+1\}$;
6. Report BR as set of broadcast agents;

Theorem 6. *Algorithm* Broadcast-on-the-Line *identifies all agents of the collection, which are able to broadcast.*

Proof. Note that $x = a_i + \overrightarrow{\mathcal{B}}(i)/2$ is the rightmost point from which the agents $1, 2, \ldots, i$ can pick up the packet and broadcast in the left direction so that the packet reaches agent 1. Also $y = a_{i+1} - \overleftarrow{\mathcal{B}}(i+1)/2$ is the leftmost point from which the agents $i+1, \ldots, n$ can pick up the packet and broadcast in the right direction so that the packet reaches agent n. Agents i and $i+1$ can communicate during this broadcast off $x \geq y$, i.e. when $\overrightarrow{\mathcal{B}}(i) + \overleftarrow{\mathcal{B}}(i+1) \geq 2(a_{i+1} - a_i)$ which is verified at line 4 of the algorithm. All such pairs of agents are included in the broadcasting set BR.

Observe that, the positive amounts of E in Convergecast-on-the-Line and $\overrightarrow{\mathcal{B}}(i) + \overleftarrow{\mathcal{B}}(i) - e(i) \geq 0$ in Broadcast-on-the-Line equal the maximal amount of energy which may be suppressed from agent i and the corresponding communication is still possible.

3 The tree environment

We suppose that the agents are placed at the nodes of the undirected tree. Observe that for each problem: the data delivery, convergecast and broadcast

the tree may be truncated so that each leaf contains an initial position of an agent. Indeed, in neither problem it makes sense to visit subtrees containing no energy source.

We developed basic ideas and tables for the line, now they can be extended to trees. The tables Δ for lines were computed locally, looking only at neighboring agents. Simiarly for trees, the value of the corresponding tables for a node in a tree is computed looking at neighbors of this node.

3.1 Data delivery in the tree

The delivery problem for a tree is easily reducible to the case of a line.

Theorem 7. *We can solve delivery problem and construct delivery-scenario on the tree in linear time.*

Proof. We find the path π connecting s with t in a given undirected tree T. Suppose we remove edges of the path. The tree splits into several subtrees *anchored* at vertices of π. For each such subtree we direct all edges towards the root, which is a vertex of π. The leaves in these trees are sending their energies towards their roots accumulating energies of intermediate nodes. This way we reduce the delivery on the tree to the delivery on the line π. This completes the proof.

3.2 Convergecast in the tree

We can reduce the convergecast and broadcast problems on general trees to trees with degree at most 3. If a vertex v has neighbors $v_1, .., v_k$ we can change it locally to small binary internal subtree with leaves $v_1, .., v_k$ by adding several edges with zero cost. Hence we assume from now that in our undirected tree each vertex has at most 3 neighbors.

Though the input tree is undirected, we can consider direction of edges. For each undirected edge (u, v) we consider two directed edges $u \to v,\ v \to u$.
Define the subtree $T(u, v)$ as the connected component containing u and resulting from T by removing the edge (u, v), see Figure 3.

For each directed edge $u \to v$ of the tree we define $\Delta(u, v)$ as the cost of moving all packets from the subtree $T(u, v)$ to its root u without interacting with any node outside $T(u, v)$. The table Δ is a generalization of the table Δ to trees.

We can use similar algorithm as computing locally values $\overrightarrow{\Delta}(v)$ on the line in the previous section, consequently we have the following fact.

Lemma 8.
(a) *Assume u, u_1, u_2 are neighbors of the vertex v. Then, knowing $\Delta(u_1, v)$, $\Delta(u_2, v)$ we can compute $\Delta(v, u)$ in constant time.*
(b) *If $\Delta(v, w),\ \Delta(w, v)$ are known then we can compute in constant time all convergecast points on the edge (v, w) as its subinterval.*

Let us root the tree T at some vertex *root* obtaining a rooted version $\overline{T}$ of T.

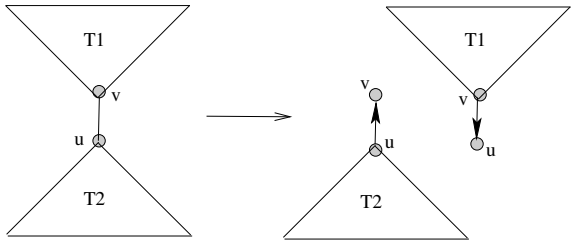

Fig. 3. Testing if there is a *convergecast point* on the undirected edge (u, v) is reduced to computation of the costs $\Delta(u, v)$ and $\Delta(v, u)$ of moving all packets in the trees $T2 = T(u, v)$ and $T1 = (v, u)$.

Observation 9. *For each node $v \in \overline{T}$, $v \neq root$, there are directed edges $v \rightarrow father(v)$, (outgoing) and $father(v) \rightarrow v$ (ingoing edge). Assume we are to mark all directed edges, but we have to obey the rule: if an edge $v \rightarrow w$ is marked then all edges ingoing into v from nodes other than w should be marked. We can obey the rule and mark all edges by first traversing the tree in postorder marking for each visted node v the edge $v \rightarrow father(v)$. Then we can traverse in preorder and for each visited node v mark the edge $father(v) \rightarrow v$.*

We can replace the operation of *marking* an edge e by the computation of table $\Delta(e)$. Using the observation above our algorithm computing convergecast points on T can be written as the following pseudocode:

ALGORITHM Convergecast on the tree T;
$root$:= any node of T; $\overline{T}$:= the directed version of T rooted at $root$;
for each node $v \neq root$ of $\overline{T}$ in *postorder* **do**
    Compute $\Delta(v, father(v))$
for each node $v \neq root$ of $\overline{T}$ in *preorder* **do**
    Compute $\Delta(father(v), v)$
for each undirected edge (v, w) of T **do**
    compute in constant time, knowing $\Delta(v, w)$, $\Delta(w, v)$
    all convergcast points on the edge (v, w) as its subinterval;
return the set of convergcast points as sets of subintervals of edges;

The next theorem follows now from Lemma 8.

Theorem 10. *The convergecast problem for undirected trees can be solved in linear time.*

3.3 Broadcasting in the tree

We use terminology from the previous section and ideas from broadcasting on line. Define $\mathcal{B}(u, v)[i, j]$ as the potential of energy at the vertex u for the delivery of initial information of i agents placed initially at u to all agents in subtree $T(u, v)$, using only agents in $T(u, v)$, assuming that at the end we have j agents at u. It is similar to the definition of $\mathcal{B}$ for broadcasting on the line. Hence each $\mathcal{B}(u, v)$ is a $n \times n$ table. The use of quadratic tables results in nonlinear algorithm.

Observation 11. *The main innovation here is introducing many agents at the same node. For example suppose we have a tree with single branch of k nodes, each with single agent, from r to u. The edges of this branch have zero cost. However v has k outgoing edges with large cost. Then optimal broadcasting of such a tree from r first gathers all k agents at node v, then each agent travels a different edge from v.*

Theorem 12. *The broadcasting on a tree can be done in polynomial time.*

Sketch of the proof We can show the following claim, analogous to Lemma 8.

Claim.
(a) Assume u, u_1, u_2 are neighbors of vertex v. Then we can compute $\mathcal{B}(v, u)$ in polynomial time, knowing tables $\mathcal{B}(u_1, v)$, $\mathcal{B}(u_2, v)$.
(b) If tables $\mathcal{B}(u, w)$ are known for each neighbor of u then we can check in polynomial time if u is a valid source of broadcast.

The tables $\mathcal{B}$ for each edge are computed locally in a similar way as values of Δ for edges, traversing first the rooted version of the tree in postorder, and then in preorder.

Each edge can be processed in polynomial time, due to the claim, hence the whole algorithm works in polynomial time. We omit the details.

4 NP-Completeness for digraphs and graphs

We use the following NP-complete problem:

Integer Set Partition: For a given set X with integer weights check if X can be partitioned into two disjoint subsets with equal total sums of weights.

Theorem 13. *The delivery problem is NP-complete for general directed graphs.*

Proof. Consider the graph of the form shown in Figure 4. The set of nodes consists of the set of middle nodes (the set X) and three additional nodes s, t, a. Initially there are agents only in middle nodes.
Denote by E the total energy of all agents. Hence $E = \alpha + \beta$.

Energy of an element $x \in X$ is set to be $e(x) = w(x)$. There is no energy in other points. The length of en edge $x \to s$ equals $w(x)/3$. The length of the edge $s \to a$ is $E/3$ and the length of the edge $a \to t$ is $E/2$. The lengths of other edges are zero.
Assume X_1, X_2 are the sets of agents which move to s and a, respectively. Let

$$\alpha = \sum_{x \in X1} w(x), \ \beta = \sum_{x \in X2} w(x)$$

Total energy coming to s is $\frac{2}{3}\alpha$ and total energy coming to a is

$$\frac{2}{3}\alpha - (\alpha + \beta)/3 + \beta = \alpha/3 + \frac{2}{3}\beta$$

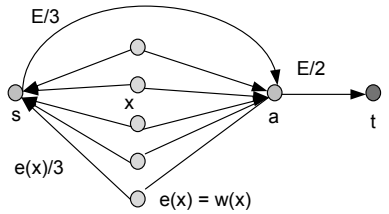

Fig. 4. The energy of all three non-middle nodes is zero. The weight $w(x)$ of each middle node x equals agent's energy $e(x)$. Delivery from s to t is possible only iff the set of weights can be partitoned into two sets of the same sum.

Delivery from s to t is possible if and only if

$$\frac{2}{3}\alpha \;\geq\; E/3 \;=\; (\alpha+\beta)/3 \quad \text{and} \quad \alpha/3+\frac{2}{3}\beta \;\geq\; E/2 \;=\; (\alpha+\beta)2$$

It is equivalent to:

$$\alpha \geq \beta \quad \text{and} \quad \beta \geq \alpha.$$

Consequently $\alpha = \beta$. In other words delivery from s to t is possible iff the integer partition problem is solvable. NP-completeness of the delivery problem follows from NP-completeness of the integer partition problem.

Theorem 14. *The delivery, convergecast and broadcast problems are NP-complete for general undirected graphs.*

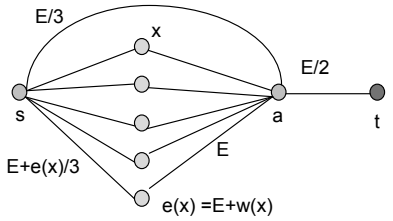

Fig. 5. The undirected version of the graph from Figure 4. The weights of middle nodes and lengths of edges incident to middle nodes are increased by E.

Proof. We take an undirected version of the graph from the previous proof, see Figure 5. Additionally for each $x \in X$ we increase by E the energy of x and edges $x \to s$, $x \to a$. The agents are placed initially in each node of the graph.
Delivery. Due to the drastic increase of the lengths of edges in the delivery from s to t the edges from X to other nodes can be used only once and in one direction (from X to outside part).
Convergecast. Now we can take t as the convergast node. The problem reduces to the delivery from s to t, since all energy is in X and each node contains an agent. Hence the problem is reduced to the delivery in the directed graph from the previous point.
Broadcast. The broadacast from s reduces to the delivery from s to t, since in the delivery all agents from X can move to s or to a. The information from s, if it has to arrive at t, should also arrive at a, where agents coming from X to a might get it.

5 Final Remarks

It is rather surprising that, without energy exchange, even the simplest problem of data delivery is NP-complete in the simplest environment of the line, while, as we have shown in this paper, all considered communication problems with energy exchange are solvable in polynomial time even for tree networks. On the other hand it is not surprising that energy exchange in general graphs does not help and the problems are NP-complete.

References

1. J. Anaya, J. Chalopin, J. Czyzowicz, A. Labourel, A. Pelc, Yann Vaxés: Collecting Information by Power-Aware Mobile Agents. *Proc. DISC* (2012), pp. 46-60.
2. S. Albers: Energy-efficient algorithms. *Comm. ACM 53(5)*, (2010), pp. 86-96.
3. S. Albers, M.R. Henzinger. Exploring unknown environments. *SIAM J. on Comput.*, 29 (4), pp.1164-1188.
4. B. Awerbuch, O. Goldreich, D. Peleg and R. Vainish: A Trade-Off between Information and Communication in Broadcast Protocols. J. ACM 37(2): 238-256 (1990).
5. V. Annamalai, S. K. S. Gupta, and L. Schwiebert, On Tree-Based Convergecasting in Wireless Sensor Networks, *Wireless Communications and Networking, IEEE*, vol. 3 (2003), pp. 1942 - 1947.
6. J. Augustine, S. Irani, C. Swamy. Optimal powerdown strategies. *SIAM J. Comput. 37* (2008), pp. 1499-1516.
7. I. Averbakh, O. Berman. A heuristic with worst-case analysis for minimax routing of two traveling salesmen on a tree. *Discr. Appl. Math.*, 68 (1996), pp. 17-32.
8. Y. Azar. On-line load balancing. In: *A. Fiat and G.Woeginger, Online Algorithms: The State of the Art, Springer* LNCS 1442, (1998), pp. 178-195.
9. R.A. Baeza-Yates, R. Schott: Parallel Searching in the Plane. *Comput. Geom. 5*, (1995), pp.143-154.
10. R. Bar-Yehuda, O. Goldreich, A. Itai: On the Time-Complexity of Broadcast in Multi-hop Radio Networks: An Exponential Gap Between Determinism and Randomization. J. Comput. Syst. Sci. 45(1): 104-126 (1992).
11. J. Chalopin, R. Jacob, M. Mihalak, P. Widmayer: Data Delivery by Energy-Constrained Mobile Agents on a Line. *Proc. ICALP (2)*, (2014), pp. 423-434.
12. M. Dynia, M. Korzeniowski, C. Schindelhauer. Power-aware collective tree exploration. In *Proc. of ARCS* (2006), pp. 341-351.
13. P. Fraigniaud, L. Gasieniec, D. Kowalski, A. Pelc. Collective tree exploration. In *Proc. LATIN*, (2004), pp. 141-151.
14. G. Frederickson, M. Hecht, C. Kim. Approximation algorithms for some routing problems. *SIAM J. on Comput.*, 7 (1978), pp. 178-193.
15. S. Irani, S.K. Shukla, R. Gupta. Algorithms for power savings. *ACM Trans. on Algorithms*, Vol. 3, No. 4, Article 41, (2007).
16. A. Kesselman and D. R. Kowalski, Fast distributed algorithm for convergecast in ad hoc geometric radio networks, *Journal of Parallel and Distributed Computing*, Vol. 66, No. 4 (2006), pp. 578-585.
17. L. Krishnamachari, D. Estrin, S. Wicker. The impact of data aggregation in wireless sensor networks. *ICDCS Workshops*, (2002), pp. 575-578.
18. R. Rajagopalan, P. K. Varshney, Data-aggregation techniques in sensor networks: a survey *Comm. Surv. and Tutorials*, Vol. 8, No. 4, (2006), pp. 48 - 63.
19. F.F. Yao, A.J. Demers, S. Shenker, A scheduling model for reduced CPU energy. In *Proc. of 36th FOCS* (1995), pp. 374-382.